

SYSTEM AND METHOD FOR INCREASING SYSTEM RESOURCE AVAILABILITY IN DATABASE MANAGEMENT SYSTEMS

Field of the Invention

5 The present invention relates generally to database management systems and more specifically to a system and method for increasing system resource availability in database management systems.

Background of the Invention

10 Transactional database systems are commonly implemented in a client-server environment in which applications access a database server to obtain data in response to requests submitted to the database server. For example, a relational database management system may be implemented on a server and applications will access the server by issuing relational database statements to the server. The database server will return results and/or error codes to the applications. In a transactional database system, records may be altered by an application, but until the application commits the changes, the altered database data will
15 generally not be accessed by other applications. Database resources may be used by an application in processing a transaction and during the transaction such resources will be therefore unavailable to other applications. When the application ends the transaction by a commit or a rollback the database system resources will be released for potential use by all applications.

20 An example of database system resources that are made unavailable during a transaction are locks. In many cases, database management systems are designed to permit multiple applications to concurrently access data. To allow such concurrent access, database management systems typically include locks that are available for acquisition by applications to ensure the correct serialization of access to data. Database applications protect changed
25 (inserted/deleted/updated) but uncommitted records by placing a lock on such changed records (in relational systems such a lock is a row lock). Such a lock prevents concurrent applications from accessing uncommitted records. An application holding one or more locks will release all the locks that the application holds when the current transaction ends with either a commit or a rollback.

According to certain database designs, each application is provided with a user configurable upper limit of on the number of locks that may be held at any one time. In addition, a system will have an implicit or explicit limit on the total number of locks held by all applications. Once this limit is reached, any attempt to lock an additional record may result in lock escalation. In lock escalation the application locks the target table (or other database structure) to protect all records and to allow the release of the previously held row locks. As a result, locks are made available in the system (typically, all locks require the same amount of system resources, so replacing multiple row locks with a single table lock frees system resources).

An online application is one that does not limit other concurrent applications from accessing committed (although not uncommitted) data in both read and write mode. An application holding a table lock is by definition offline, because concurrent applications are prevented from accessing the table. Therefore, for an application to stay online the application must be able to avoid lock escalation. For this reason, it is desirable to ensure that applications do not reach their defined limit of held locks.

Another system resource that can become a limiting factor in a database system that supports multiple concurrent online applications is the active log space. A characteristic of certain database management systems is that certain changes to data (for example, a change made in a relational database using SQL statements, as opposed to some data movement utilities) will result in log records being written prior to updating the database. Such logs are used in the process of database recovery in which the database is returned to a consistent state after a failure (crash recovery, or a roll-forward recovery following a database restore from a backup). If the space needed to log the change is not available in the database logs, the database language statement will fail and in most cases the current transaction will be rolled back. Active log space consumed by an application cannot be freed until the active transaction ends (commit or rollback). Since log space is shared between all applications using the same database, this resource may become scarce when multiple applications are running concurrently.

To increase the availability of resources such as record locks and log space to applications in a database environment, an autocommit protocol may be adopted. In such an

approach, where possible, an application will commit the current transaction after every database system statement is issued (for example after the each SQL statement is issued in a relational database system). This provides for the release to other applications of system resources that are held by the application until the end of a transaction. After each commit, the resources, such as locks and the active log space, are made available in the database system. However, there is a cost incurred resulting from performing commits based only on the issuance of database system statements and not based on the actual or predicted demand for system resources. This may be particularly inefficient in distributed environments (multiple physical machines running the same instance of a database manager) when multiple requests might have to be executed on different database nodes.

It is therefore desirable to provide a method and system for applications to potentially release system resources when a server determines that system resources have become scarce.

Summary of the Invention

Accordingly, the system and method of the present invention provide for improved support for increasing system resource availability in database management systems.

According to an aspect of the invention a method is provided for optimizing the number of commits performed with the objective of making system resources available to applications. According to an aspect of the invention, a method is provided for keeping concurrent applications online by triggering commits in a manner so as to make available shared system resources that are freed on termination of transactions.

According to another aspect of the invention there is provided, for a database server, a computer program product for directing the database server to selectively release database system resources associated with a database system operatively coupled with the database server, the database server accessible by applications, the computer program product including a computer usable medium having computer readable program code means embodied in said medium, the computer readable program code means including computer readable program code means for receiving requests from one or more applications, computer readable program code means for determining whether a specified application is submitting requests using an online protocol, and computer readable program code means for returning defined error

condition indicators to the specified application when the application is determined to be using the online protocol and when identified system resources are determined to be scarce, whereby the specified application issues a commit when it is using the online protocol and one of the defined error condition indicators is returned to the application.

5 According to another aspect of the invention there is provided a computer-implemented method for directing a database server to selectively release database system resources associated with a database system operatively coupled with the database server, the database server accessible by applications, the method including the steps of receiving an online
10 protocol indicator from an application, returning a defined error condition indicator to the application when the application has forwarded the online protocol indicator and when identified system resources are determined to be scarce, and receiving a commit statement from the application when a response from the database server is determined to comprise the defined error condition indicator.

15 According to another aspect of the invention there is provided a computer program product for directing an application to access a database server, the database server adapted to selectively release database system resources associated with a database system operatively coupled with the database server, the database server accessible by applications, the computer
20 program product including a computer usable medium having computer readable program code means embodied in said medium, the computer readable program code including computer readable program code means for submitting requests to the database server, computer readable program code means for indicating that requests are submitted using an online protocol, computer readable program code means for monitoring responses from the database server for defined error condition indicators, when requests have been submitted
25 using the online protocol, and computer readable program code means for issuing a commit statement to the database server when a response from the server is determined to be a defined error condition indicator.

The present invention thus permits database management systems to increase the availability of system resources to applications.

Brief Description of the Drawings

In the drawings which illustrate by way of example only a preferred embodiment of the invention,

Figure 1 is a block diagram illustrating a schematic representation of an application and server in accordance with the preferred embodiment;

5 Figure 2 is a flow chart showing the operation of the preferred embodiment on the server of the preferred embodiment;

Figure 3 is a flow chart showing the operation of an application in accordance with the preferred embodiment.

10 **Detailed Description of the Invention**

The preferred embodiment of the invention is described with reference to a transactional database system and which is an SQL (Structured Query Language) relational database system. It will be appreciated by those skilled in the art that the invention may be implemented in transactional databases other than that used to illustrate the preferred
15 embodiment.

The preferred embodiment is defined in terms of a database server accessed by applications. A database server is a computer program that manages data in a database and ensures its consistency. The term database server is synonymous with the term "database manager". An application is a computer program that accesses data stored in a database managed by a database server.
20

The database server of the preferred embodiment processes transactions. A transaction is a set of one or more database statements (in the preferred embodiment, SQL statements) that execute as a single operation. The term transaction is synonymous with the term unit of work. Transactions are defined by applications and the manner in which
25 applications send database statements to the database server. In the description of the preferred embodiment, a multi-transactional application is an application that executes a sequence of transactions and that does not impose strictly defined transaction boundaries.

As will be appreciated by those skilled in the art, a database application may permit, or alternatively may limit, other concurrent applications from accessing (in both read and write mode) committed data maintained in the database. Where an application permits other concurrent applications to have such access to committed data, the application is considered to be online. To allow access to data in the database server it is preferred to keep all applications online.

A commit in the database system of the preferred embodiment terminates a transaction and commits the database changes that were made by an application in that transaction. Database changes are said to be committed if the transaction that made them was successfully committed. Otherwise, database changes are said to be uncommitted. An alternative to a commit is a rollback. A rollback terminates a transaction and backs out the database changes that were made by that transaction.

In the preferred embodiment, database statements executed within a single transaction are either committed or rolled back. Hence a transaction can be thought of as a single logical operation. This property of a transaction is referred to as atomicity. For an application to be considered to be multi-transactional, as the term is used in this description, the application must be able to permit a transaction to be interrupted by a commit without affecting the logic of the originally defined transaction. In other words, any given statement in a transaction must be capable of being placed in a different unit of work, without affecting the logic of the application for the application to be considered multi-transactional. Any application guaranteeing atomicity of a sequence of operations (i.e. the whole sequence has to be either committed or rolled back) is not multi-transactional.

In the preferred embodiment, locks are made available by the database system to provide a means of ensuring data integrity while attempting to reduce resource conflicts. Each application using a shared resource will acquire a lock on the resource to ensure its consistency throughout the transaction. It will be understood that a lock is, itself, a system resource.

Figure 1 is a block diagram illustrating an extremely simple example arrangement of a database server and applications in accordance with the preferred embodiment. Figure 1 shows applications 8, 10. Both applications are able to access database server 12. Server 12

includes tables 14 logs 16, and lock list 17. In SQL database server 12 in the preferred embodiment, relational data is stored in tables 14.

5 In use, applications 8, 10 submit database requests to server 12. In response, database server 12 sends results and/or error codes to the appropriate application 8, 10. As transactions are processed, their execution is tracked within the active log files 16. Information in log files 16 is used during a transaction rollback and during a database recovery. Lock list 17 maintains a list of locks available and used in the database system.

10 Database server 12 supports applications 8, 10 accessing database server 12 using a specified online protocol. Where one of applications 8, 10 accesses database server 12 in using the online protocol, database server 12 follows a sequence of operations (as set out in more detail below) that will seek to increase the system resources available to applications 8, 10. It will be appreciated that other implementations may identify or specify applications as being online in different manners but in the preferred embodiment an online flag is used to specify whether a given application is accessing the database server in the online mode.

15 When one of applications 8, 10 sets online flag 18 in database server 12 it means that the application is (for at least the time that online flag 12 is set) a multi-transactional application. In the terminology of the preferred embodiment, the application is using the online protocol. The application may issue a commit at any time that online flag 18 is set without changing the logic of the application. As will be appreciated, database server 12 needs
20 to be informed as to whether the requestor application is using the online protocol. This can be done either on the application level, or on the statement level.

25 An application level online flag can be set (or reset) in the memory of database server 12 by executing a new database request. In the preferred embodiment, this request can be implemented either as a new SQL statement, or as a stored procedure executing on server 12. Alternatively, or additionally, a statement level online indicator is sent to the server as a part of each SQL statement. In the preferred embodiment, this is achieved by extending the SQL syntax to include a new keyword (such as 'ONLINE') that is able to be specified for various SQL statements. Upon recognizing the specified keyword, a compiler in server 12 (not shown) setups the necessary information in database server 12 before the request containing the

statement with the keyword is executed at server 12.

In the preferred embodiment, both approaches may be implemented. Server 12 will execute the operations as set out below defined for the online protocol if either online flag 18 is set by an application request or the statement being executed carries the online indicator. A statement level online indicator offers more flexibility, but the cost is higher because the information has to be distributed (potentially to many database nodes in a distributed system) with every SQL request. Setting the single application level flag in server 12 using a request is more efficient because only a one-time cost is incurred. As will be appreciated, online flag 18 may be implemented in a number of different ways and the existence of a single data element capable of storing different values is not necessary to achieve the functionality of online flag 18.

According to the preferred embodiment, when execution of a request from an application occurs under the online protocol, the server will provide a specified error code to the requesting application when system resources are identified as being scarce. Figure 2 is a flow chart that illustrates the operation of server 12 in accordance with the preferred embodiment in which active log and lock list resources are subject to the online protocol. In the example shown, server 12, when executing in accordance with the online protocol, will return the specific error code to the application if either the lock list is full, or if the active log space is full.

Server 12 receives a request from one of applications 8, 10 (box 20 in Figure 2). In response to the received request, server 12 performs SQL processing (box 22). Server 12 will interrogate logs 16 to determine if the active log is full (decision block 24). If the active log is not full server 12 will determine if there are available locks (decision block 26). As will be appreciated, the definition of “full” in the above description may be dependent on system design and may be determined such that there remains some portion of the active log or lock list that is available when the “full” condition is met.

If neither the active list is full nor the lock list is full, server 12 will continue with processing (block 28). When the processing is complete server 12 will return a result or an error code to the database application (block 30). As will be appreciated, server 12 may

interrogate the active logs and the lock list more than a single time.

If the active log is determined to be full (decision block 24) then server 12 determines whether the online protocol is being followed (decision block 32). If online protocol flag 18 is not set, then server 12 returns to the requesting application. Alternatively, if online protocol
5 flag 18 is set, server 12 returns the log full error to the requesting application (blocks 34, 30).

Similarly, if the lock list is determined to be full (decision block 26) then the determination is made as to whether the online protocol applies (decision block 36). If online protocol flag 18 is set, the a lock full error code is returned by server 12 to the requesting application (blocks 34, 30). If online flag 18 is not set, server 12 escalates locks (block 38) and
10 more processing is carried out (block 28). In this alternative, the requesting application effectively goes offline.

As may be seen, when an application is in online protocol mode and one of the system resources reaches a defined level of scarcity, server 12 returns an error code to the requesting application. The code is related to, and in the preferred embodiment the code specifies, the
15 particular resource that has become scarce.

From the application side of the preferred embodiment, to specify use of the online protocol either a one time request is executed against database server 12, or the specified online keyword is included when preparing every SQL statement, depending on whether the application level or statement level flag specification is more appropriate for the application.

While processing SQL statements, each of applications 8, 10 monitors the error codes potentially returned from server 12. By setting online protocol flag 18 an application indicates that it is a multi-transactional application and that consequently executing a commit while the flag is set will not affect the logic of the application. Thus, when a specific error code relating to system resource scarcity is received by the application from database server 12, application
20 is able to trigger a commit and therefore release held resources, without affecting the application logic. In the example of the preferred embodiment, executing a commit will potentially free lock space or active log space resources. If a commit fails to free the resources (due to an error), a re-executed statement will fail with the same specific error code. In such
25 a case, the transaction will then be rolled back.

This process of releasing resources is shown in the flow chart of Figure 3. The application process start is shown in block 50. If an application level online flag is used then a request to the server to setup the appropriate flag is made (blocks 52, 54). If a statement level online indicator is used, then an SQL online statement is prepared (blocks 56, 58). If there is no statement level online indicator used, then the application prepares the SQL statement without the indicator (block 60). The prepared SQL statement is then executed (block 62). The application monitors the responses from the server and if an error is returned it is checked to see the type of error (decision blocks 64, 66). A commit is executed if the error returned is a log full error or a lock list full error (blocks 66, 68). After a commit, the statement is re-executed (block 70). The return from the server is then checked and if there is an SQL error a rollback is executed (blocks 72, 76). Otherwise, the application determines whether there are more SQL statements to process. If there are then the application returns to block 56. Otherwise a commit is executed to terminate the defined transaction (blocks 74, 78). The database application process then stops (block 80).

As may be seen from the above description, system resources are potentially made available to applications by applications specifying an online protocol to the database server. The above example refers to lock and log resources but other resources that can be freed by committing database changes may be similarly treated. By the server providing specific error codes that relate to the scarcity of database system resources, the application is able to issue a commit to allow resources to be freed up and to prevent, for example, lock escalation.

Various embodiments of the present invention having been thus described in detail by way of example, it will be apparent to those skilled in the art that variations and modifications may be made without departing from the invention. The invention includes all such variations and modifications as fall within the scope of the appended claims.